

PATENT APPLICATION
DOCKET NO.:200209000-1

"EXPRESS MAIL" Mailing Label No..EV331251302US.....
Date of Deposit.....AUGUST 6, 2003.....

DATA SELECTION CIRCUIT FOR PERFORMANCE COUNTER

PRIORITY UNDER 35 U.S.C. §119(e) & 37 C.F.R. §1.78

[0001] This nonprovisional application claims priority based upon the following prior United States provisional patent application entitled: "General Purpose Counters for Performance, Debug and Coverage," Application No.: 60/469,180, filed May 9, 2003, in the name(s) of Richard W. Adkisson and Tyler J. Johnson, which is hereby incorporated by reference.

CROSS-REFERENCE TO RELATED APPLICATIONS

[0002] This application is related to U.S. Patent Application Serial No. _____, filed _____ entitled COVERAGE CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200208996-2); U.S. Patent Application Serial No. _____, filed _____ entitled COVERAGE DECODER CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200208997-1); U.S. Patent Application Serial No. _____, filed _____ entitled GENERAL PURPOSE PERFORMANCE COUNTER (Docket No. 200208999-1); U.S. Patent Application Serial No. _____, filed _____ entitled ZEROING CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200209001-1); U.S. Patent Application Serial No. _____, filed _____

entitled MATCH CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200209002-1); and U.S. Patent Application Serial No. _____, filed _____ entitled INCREMENT/DECREMENT CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200208998-1), all of which are hereby incorporated by reference in their entirety.

[0003] Related subject matter disclosed in the following commonly owned co-pending U.S. patent applications: (i) A BUS INTERFACE MODULE, filed March 28, 2003; Application No. 10/402,092; and (ii) AN INTEGRATED CIRCUIT, filed March 28, 2003; Application No. 10/402,034, is hereby incorporated by reference.

BACKGROUND

[0004] Increasing demand for computer system scalability (i.e., consistent price and performance and higher processor counts) combined with increases in performance of individual components continues to drive systems manufacturers to optimize core system architectures. One such systems manufacturer has introduced a server system that meets these demands for scalability with a family of application specific integrated circuits ("ASICs") that provide scalability to tens or hundreds of processors, while maintaining a high degree of performance, reliability, and efficiency. The key ASIC in this system architecture is a cell controller ("CC"), which is a processor-I/O-memory interconnect and is responsible for communications and data transfers, cache coherency, and for providing an interface to other hierarchies of the memory subsystem.

[0005] In general, the CC comprises several major functional units, including one or more processor interfaces, memory units, I/O controllers, and external crossbar interfaces all interconnected via a central data path ("CDP"). Internal signals from these units are collected on a performance monitor bus ("PMB"). One or more specialized performance counters, or performance monitors, are connected to the PMB and are useful in collecting data from the PMB for use in debugging and assessing the performance of the system of which the CC is a part. Currently, each of the performance counters is capable of collecting data from only one preselected portion of the PMB, such that the combination of all of the performance counters together can collect all of the data on the PMB. While this arrangement is useful in some situations, there are many situations in which it would be advantageous for more than one of the performance counters to access data from the same portion of the PMB. Additionally, it would be advantageous to be able to use the performance counters in the area of determining test coverage. These applications are not supported by the state-of-the-art performance counters.

SUMMARY

[0006] In one embodiment, the invention is directed to a data selection circuit for a general purpose performance counter ("GPPC") connected to a bus carrying debug data. The data selection circuit is connected to receive the debug data and comprises logic for receiving the debug data as a plurality of N -bit portions of block-aligned data and outputting a designated one of the N -bit portions; and

circuitry for providing to the receiving logic a control signal for designating one of the N -bit portions.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a block diagram illustrating general purpose data collection in a logic design;

[0008] FIG. 2 is a block diagram of a general purpose performance counter according to one embodiment;

[0009] FIG. 3 is a more detailed block diagram of the general purpose performance counter of FIG. 2;

[0010] FIG. 4 illustrates a method in which signals are mapped from an observability bus to a performance counter in accordance with one embodiment; and

[0011] FIG. 5 is a more detailed block diagram of a MUX circuit of an sm_sel circuit of the general purpose performance counter of FIG. 3.

DETAILED DESCRIPTION OF THE DRAWINGS

[0012] In the drawings, like or similar elements are designated with identical reference numerals throughout the several views thereof, and the various elements depicted are not necessarily drawn to scale.

[0013] FIG. 1 is a block diagram of general purpose data collection in a logic design. As shown in FIG. 1, the state space 100 of a logic design under consideration is driven to data collection and selection logic 102. The logic 102 drives a D -bit data collection, or observability, bus 104 carrying a D -bit debug_bus signal to a plurality of performance counters 106(1)-106(M). Details of one

embodiment of the logic 102 and bus 104 are provided in U.S. Patent Application Serial No. 10/402,092, filed March 28, 2003, entitled A BUS INTERFACE MODULE (Docket No. 200208674-1); and U.S. Patent Application Serial No. 10/402,034, filed March 28, 2003, entitled AN INTEGRATED CIRCUIT (Docket No. 200209004-1), each of which is hereby incorporated by reference in its entirety.

[0014] In one embodiment, D is equal to 80, M is equal to 12, and performance counters 106(1)-106(M -1) are general purpose performance counters, while the remaining performance counter 106(M) increments on every clock cycle. As will be illustrated below, the general purpose performance counters are "general purpose" in that each of them is capable of accessing any bit of the 80-bits on the bus 104; moreover, all of them may access the same block of bits and do the same or different performance calculations thereon.

[0015] FIG. 2 is a block diagram of a general purpose performance counter 200, which is identical in all respects to each of the performance counters 106(1)-106(M -1) (FIG. 1), in accordance with one embodiment. As will be described in greater detail below, the performance counter 200 can be used to perform general purpose operations to extract performance, debug, or coverage information with respect to any system under test (SUT) such as, for instance, the system state space 100 shown in FIG. 1. The performance counter 200 includes an AND/OR circuit 201, a match/threshold circuit 202, an sm_sel circuit 204, an szero circuit 206, and a counter circuit 208.

[0016] In general, the AND/OR circuit 201 enables access to all of the bits of the debug_bus signal coming into the

performance counter 200 via the observability bus 104. In one embodiment, as illustrated in FIGS. 2 and 3, debug_bus is an 80 bit signal. When the AND/OR circuit 201 is operating in AND mode, the circuit activates an "inc" signal to the counter circuit 208 if all of the bits of the debug_bus signal plus two bits that are appended thereto, as will be described in greater detail below, that are of interest (as indicated by the value of an 80-bit "mask" plus two bits that are appended thereto) are set. When the AND/OR circuit 201 is operating in OR mode, the circuit activates the inc signal to the counter circuit 208 if any one or more of the bits of the debug_bus signal plus the two additional bits that are of interest (as indicated by the value the mask plus the two additional bits) are set.

[0017] When the match/threshold circuit 202 is operating in "match" mode, a match portion 300 (FIG. 3) of the circuit activates a match_thresh_event signal to the AND/OR circuit 201 when an N -bit portion of the debug_bus signal selected as described in greater detail below with reference to the sm_sel circuit 204 and the szero circuit 206 matches an N -bit threshold for all bits selected by a match mask ("mmask"). In particular, for all bits of the selected N -bit debug bus signal portion that are "don't cares", the corresponding bit of mmask will be set to 0; conversely, for all bits of the selected N -bit debug bus signal portion that are not "don't cares", the corresponding bit of mmask will be set to 1. The match_thresh_event signal is one of the two bits appended to the debug_bus signal. In the illustrated embodiment, N is equal to 16.

[0018] When the match/threshold circuit 202 is operating in "threshold" mode, a threshold portion 302 (FIG. 3) of the circuit 202 activates the match_thresh_event signal to the AND/OR circuit 201 when an S -bit portion of the debug_bus signal selected and zeroed as described in greater detail below with reference to the sm_sel circuit 204 and the szero circuit 206 is equal to or greater than the threshold. In the illustrated embodiment, S is equal to $N/2$, or 8.

[0019] Additional details regarding operation of the match/threshold circuit 202 are provided in U.S. Patent Application Serial No. _____, filed _____ entitled MATCH CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200209002-1).

[0020] The sm_sel circuit 204 selects an N -bit portion of the debug_bus signal aligned on a selected 10-bit block boundary into both the match portion 300 and the threshold portion 302 (FIG. 3) of the match/threshold circuit 202 and to a sum input of the counter circuit 208. As previously stated, in the illustrated embodiment, N is equal to 16. The szero circuit 206 zeroes out none or all but one of S bits aligned on a selected 10-bit block boundary into the threshold portion 302 of the match/threshold circuit 202 and the sum input of the counter circuit 208. In the illustrated embodiment, S is equal to eight. The selected 10-bit block boundary is identified by the value of a three-bit control signal sm_sel input to the sm_sel circuit 204.

[0021] In one embodiment, the sm_sel circuit 204 comprises a MUX circuit comprising N eight-input MUXes, which for the sake of simplicity are collectively represented in FIG. 3 by a single eight-input MUX 303, the output of which is selected

with an sm_sel[2:0] signal. The sm_sel[2:0] signal, which may be provided from a selection control circuit 305 (such as a control status register ("CSR"), for example) selects which of the eight 10-bit-block-aligned N-bit portions of the debug_bus signal will be output to the match/threshold circuit 202 and to the szero circuit 206. This concept will be illustrated in greater detail with reference to FIG. 4.

[0022] Additional details regarding the operation of the szero circuit 206 are provided in U.S. Patent Application Serial No. _____, filed _____ entitled ZEROING CIRCUIT FOR PERFORMANCE COUNTER (Docket No. 200209001-1).

[0023] In one embodiment, each general purpose performance counter, such as the performance counter 200, is 48 bits plus overflow. The performance counter 200 is general purpose in that it looks at all D bits of the debug_bus signal for an event mask plus two extra events, eight separate selections of 16 bits for the match compare operation and eight separate selections of eight bits for the threshold compare and the accumulate operations. The eight bits for the threshold compare and the accumulate operations are the bottom eight bits of the 16 bits selected for the match compare operation. Those 16 bits are aligned to 10 slot boundaries as shown in an exemplary mapping arrangement illustrated in FIG. 4.

[0024] In FIG. 4, an events signal 400 comprises the debug_bus signal, designated in FIG. 4 by reference numeral 401, the match_threshold_event signal, designated by reference numeral 402 and a logic 1 bit, designated by reference numeral 404. The debug_bus signal 401 comprises bits [79:0] of the events signal 400; the match_threshold_event signal 402 comprises bit [80] of the

events signal, and the logic 1 bit 404 comprises bit [81] of the events signal.

[0025] As best illustrated in FIG. 3, the events signal 400 (i.e., the debug_bus signal with the match_threshold_event signal and the logic 1 appended thereto) are input to a first logic stage 304 of the AND/OR circuit 201 for purposes that will be described in greater detail below.

[0026] Referring again to FIG. 4, a composite mask signal 410 comprises an 80-bit mask signal, designated by a reference numeral 412, a match_threshold_event mask ("TM") bit, designated by reference numeral 414, and an accumulate ("acc") bit, designated by reference numeral 416. The mask signal 412 comprises bits [79:0] of the composite mask signal 410; the TM bit 412 comprises bit [80] of the composite mask signal, and the acc bit 416 comprises bit [81] of the composite mask signal. As best illustrated in FIG. 3, each bit of the composite mask 410 (i.e., the mask signal with the TM and acc bits appended thereto) is input to the first logic stage 304 of the AND/OR circuit 201 for purposes that will be described in greater detail below.

[0027] Continuing to refer to FIG. 4, eight 10-bit block-aligned 16-bit match selections are respectively designated by reference numerals 420(0)-420(7). In particular, the selection 420(0) comprises bits [0:15]; the selection 420(1) comprises bits [10:25]; the selection 420(2) comprises bits [20:35]; the selection 420(3) comprises bits [30:45]; the selection 420(4) comprises bits [40:55]; the selection 420(5) comprises bits [50:65]; the selection 420(6) comprises bits

[60:75]; and the selection 420(7) comprises bits [70:5] (bits above 79 wrap back to zero).

[0028] Referring again to FIG. 3 and the description of the sm_sel circuit 204 set forth above, it will be noted that an sm_sel[2:0] value of 0 will result the selection 420(0) being output from the MUX circuit represented by the MUX 303; an sm_sel[2:0] value of 1 will result the selection 420(1) being output from the MUX circuit represented by the MUX 303; an sm_sel[2:0] value of 2 will result the selection 420(2) being output from the MUX circuit represented by the MUX 303; and so on.

[0029] FIG. 5 is a more detailed block diagram of the MUX circuit represented in FIG. 3 by the MUX 303. As illustrated in FIG. 5, the MUX circuit comprises N , in this case, sixteen, eight-input MUXes 500(0)-500(15). The sm_sel[2:0] signal is provided to control inputs of each of the MUXes 500(0)-500(15). The least significant bit ("LSB") (i.e., bit 0) of each of the selections 420(0)-420(7) is input to a corresponding input of the MUX 500(0). In particular, the LSB of the selection 420(0) is input to input 0 of the MUX 500(0), the LSB of the selection 420(1) is input to input 1 of the MUX 500(0), the LSB of the selection 420(2) is input to input 2 of the MUX 500(0), and so on.

[0030] Similarly, the second LSB (i.e., bit 1) of each of the selections 420(0)-420(7) is input to a corresponding input of the MUX 500(1). In particular, the second LSB of the selection 420(0) is input to input 0 of the MUX 500(1), the second LSB of the selection 420(1) is input to input 1 of the MUX 500(1), the second LSB of the selection 420(2) is input to input 2 of the MUX 500(2), and so on.

[0031] In the same manner, the third LSB (i.e., bit 2) of each of the selections 420(0)-420(7) is input to a corresponding input of the MUX 500(2), and so on for each bit of each of the selections 420(0)-(7) and each of the remaining MUXes 500(3)-500(15) such that the most significant bit ("MSB") of the selection 420(0) is input to input 0 of the MUX 500(15), the MSB of the selection 420(1) is input to input 1 of the MUX 500(15), the MSB of the selection 420(2) is input to input 2 of the MUX 500(15), and so on.

[0032] In this manner, in one embodiment, when the sm_sel[2:0] signal is [000], all of the bits of the first selection 420(0) will be output from the MUXes 500(0)-500(15), with the LSB being output from the MUX 500(0) and the MSB being output from the MUX 500(15).

[0033] Referring back to FIG. 3, the first logic stage 304 comprises an AND portion, represented by an AND gate 304a, for bit-wise ANDing the events signal 400 with the composite mask signal 410, and an OR portion, represented by an OR gate 304b, for bit-wise ORing the inverse of the composite mask signal 410 with the events signal 400. It will be recognized that, although represented in FIG. 3 as a single two-input AND gate 304a, the AND portion of the first logic stage 304 actually comprises 82 two-input AND gates. Similarly, the OR portion of the first logic stage 304 comprises 82 two-input OR gates identical to the OR gate 304b.

[0034] The outputs of the AND portion of the first logic stage 304 are input to an 82-input OR gate 306, the output of which is input to one input of a two-input MUX 308 as an "or_result". Similarly, the outputs of the OR portion of the

first logic stage 304 are input to an 82-input AND gate 310, the output of which is input to the other input of the MUX 308 as an "and_result". A control signal ("and/or#"), which may originate from a CSR (not shown), controls whether the AND/OR circuit 201 functions in AND mode, in which case the and_result is output from the MUX 308 as the inc signal, or in OR mode, in which case the or_result is output from the MUX as the inc signal.

[0035] As a result, when the AND/OR circuit 201 is operating in the AND mode, the inc signal comprises the and_result signal and will be activated when all of the bits of the events signal 400 that are of interest as specified by the composite mask 410 are set. When the AND/OR circuit 201 is operating in OR mode, the inc signal comprises the or_result signal and will be activated when any one of the bits of the events signal 400 that are of interest as specified by the composite mask 410 is set.

[0036] The acc bit 416 of the composite mask 410 is CSR-settable. Setting the TM bit 414 in the composite mask 410 designates the match_thresh_event signal in the events signal as a bit of interest; not setting the TM bit in the composite mask will cause the value of the match_thresh_event signal in the events signal 400, and hence the result of any match or threshold operation performed by the match/threshold circuit 202, to be ignored.

[0037] Continuing to refer to FIG. 3, the operation of an embodiment of the counter circuit 208 will be described in greater detail. The counter circuit 208 is an X bit counter that can hold, increment by one, add S bits, clear, or load a value into a count value register 312. Other processing may

also occur in order to read the value of the register 312. In the embodiment illustrated in FIG. 3, X is equal to 48. Counter circuit 208 operation is enabled by setting a counter enable signal B, which comprises one input of a two-input AND gate 314. The other input of the AND gate 314 is connected to receive the inc signal from the AND/OR circuit 201. Accordingly, when the counter circuit 208 is enabled and the inc signal is activated, a logic one is output from the AND gate 314. In any other case, the output of the AND gate 314 will be a logic zero. The output of the AND gate 314 is replicated by an 8x replicator 316 and the resulting 8-bit signal is bit-wise ANDed with an 8-bit signal output from a MUX circuit 318. The inputs to the MUX circuit 318 are the sum[7:0] signal output from the szero circuit 206 and an 8-bit signal the value of which is [00000001]. The sum[7:0] signal will be output from the MUX circuit 318 when the acc signal is activated; otherwise, the [00000001] signal will be output from the MUX circuit.

[0038] An AND circuit, represented by an AND gate 320, bit-wise ANDs the signals output from the replicator 316 and from the MUX circuit 318. The resulting 8-bit signal is input to a register 322. An adder 324 adds the 8-bit signal stored in the register 322 to the 48-bit sum stored in the count value register 312. The new sum output from the adder 324 is input to a MUX circuit 326. Two other sets of inputs to the MUX circuit 326 are connected to a logic zero and a `csr_write_value`, respectively. When a `csr_write` enable signal to the MUX circuit 326 is activated, the value of `csr_write_value` is output from the MUX circuit 326 and written to the count value register 312. In this manner, a value can be loaded into the count value register 312.

Similarly, when the clear_counter signal is asserted, 48 zero bits are output from the MUX circuit 326 to the count value register 312, thereby clearing the register.

[0039] If neither the csr_write signal nor the clear_counter signal is asserted and the acc signal is asserted, the output of the adder 324 is written to the count value register 312, thereby effectively adding S bits (i.e., the value of the sum[7:0] signal) to the previous value of the count value register 312. Not enabling the counter circuit 208 results in the count value register 312 being held at its current value. Finally, to increment the value of the count value register 312 by one, the counter circuit 208 must be enabled, the inc signal must be asserted, and the acc signal must not be asserted.

[0040] FIG. 4 illustrates that the entire data collection bus 104 (FIG. 1) is available for all of the performance counters, each being represented by the performance counter 200, making them general purpose. All D bits of the debug_bus signal can be used by the AND/OR circuit 201. N bits aligned on block boundaries can be selected by the sm_sel circuit 206, enabling full coverage of the observability bus 104.

[0041] As has been described in detail hereinabove, in one embodiment, the sm_sel circuit 204 enables the selection of overlapping sections of the data collection bus 104 (FIG. 1) into more specialized and smaller computational units. In particular, the sm_sel circuit 204 of one embodiment selects N bits aligned on block boundaries into the threshold portion 302, the match portion 300, and the sum input of the counter circuit 208 (via the szero circuit 206). This enables, for

example, the match portion 300 to cover the entire range of the data collection bus 104 with overlap. As a result, the illustrated embodiment enables every general purpose performance counter in which it is implemented to select any N bits aligned on a block boundary into the match portion 300 or the threshold portion 302 of the match/threshold circuit 202, as well as to the sum input of the counter circuit 208. This general purpose selection will support, for example, a queue depth field on the data collection bus 104 to be fed to every general purpose performance counter on the bus and a P -bar histogram (where P is the number of general purpose performance counters) to be constructed by having each of P performance counters match or threshold on a different value of the queue depth field.

[0042] As previously mentioned, prior art performance counter designs were not general purpose, in that they have limited range and are designed solely for performance calculations and debug of a system design. The embodiments described herein are general purpose, in that the AND/OR circuit can perform calculations on the entire range of the data collection bus 104. The embodiments also incorporate the concept of coverage. In particular, by observing specific states in a logic design, the designer can determine how much of the state space thereof is being covered by the test vectors of a test suite. The designer can thereby gauge whether more tests need to be run and what needs to be added to fully test the entire design.

[0043] An implementation of the invention described herein thus provides a general purpose performance counter. The embodiments shown and described have been characterized as

being illustrative only; it should therefore be readily understood that various changes and modifications could be made therein without departing from the scope of the present invention as set forth in the following claims. For example, while the embodiments are described with reference to an ASIC, it will be appreciated that the embodiments may be implemented in other types of ICs, such as custom chipsets, Field Programmable Gate Arrays ("FPGAs"), programmable logic devices ("PLDs"), generic array logic ("GAL") modules, and the like. Furthermore, while the embodiments shown are implemented using CSRs, it will be appreciated that control signals may also be applied in a variety of other manners, including, for example, directly or may be applied via scan registers or Model Specific Registers ("MSRs"). Additionally, although specific bit field sizes have been illustrated with reference to the embodiments described, e.g., 16-bit threshold for pattern matching (where the bottom 8 bits are used for the threshold), 80-bit mask signal, 3-bit sm_sel, et cetera, various other implementations can also be had.

[0044] Accordingly, all such modifications, extensions, variations, amendments, additions, deletions, combinations, and the like are deemed to be within the ambit of the present invention whose scope is defined solely by the claims set forth hereinbelow.